

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	12011	delay\$3 with lock\$3 with loop\$3	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/14 14:43
L2	10508	delay\$3 with lock\$3 with circuit\$3	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/14 14:43
L3	15492	1 or 2	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/14 14:44
L4	2045636	mode	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/14 14:44
L5	5823425	power or standby or active or sleep or inactive	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/14 14:45
L6	174289	4 near5 5	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/14 14:45
L7	1630	3 and 6	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/14 14:46

EAST Search History

L8	1216	7 and memory	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/14 14:46
L9	1170	8 and clock	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/14 14:49
L10	8074	3 same clock	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/14 14:49
L11	1078	10 and 7	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/14 14:49
L12	2680	3 same memory	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/14 14:49
L13	557	12 and 7	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/14 14:49
<i>read</i> L14	506	11 and 13	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/14 14:49

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	1175	(delay\$3 with lock\$3 with loop\$3).clm.	US-PGPUB	OR	ON	2007/07/14 14:57
L2	0	(delay\$3 with lock\$3 with circuit\$3).clm.	US-PGPUB	OR	ON	2007/07/14 14:57
L3	803	(delay\$3 with lock\$3 with circuit\$3).clm.	US-PGPUB	OR	ON	2007/07/14 14:57
L4	1247	1 or 3	US-PGPUB	OR	ON	2007/07/14 14:57
L5	9384	((power or standby or active or sleep or inactive) near5 mode).clm.	US-PGPUB	OR	ON	2007/07/14 14:59
read (L6)	75	4 and 5	US-PGPUB	OR	ON	2007/07/14 14:59

? show files

File 15:ABI/Inform(R) 1971-2007/Jul 14
(c) 2007 ProQuest Info&Learning
File 16:Gale Group PROMT(R) 1990-2007/Jul 13
(c) 2007 The Gale Group
File 148:Gale Group Trade & Industry DB 1976-2007/Jul 11
(c)2007 The Gale Group
File 160:Gale Group PROMT(R) 1972-1989
(c) 1999 The Gale Group
File 275:Gale Group Computer DB(TM) 1983-2007/Jul 11
(c) 2007 The Gale Group
File 621:Gale Group New Prod.Annou.(R) 1985-2007/Jul 11
(c) 2007 The Gale Group
File 9:Business & Industry(R) Jul/1994-2007/Jul 10
(c) 2007 The Gale Group
File 20:Dialog Global Reporter 1997-2007/Jul 14
(c) 2007 Dialog
File 476:Financial Times Fulltext 1982-2007/Jul 14
(c) 2007 Financial Times Ltd
File 610:Business Wire 1999-2007/Jul 14
(c) 2007 Business Wire.
File 613:PR Newswire 1999-2007/Jul 14
(c) 2007 PR Newswire Association Inc
File 624:McGraw-Hill Publications 1985-2007/Jul 13
(c) 2007 McGraw-Hill Co. Inc
File 634:San Jose Mercury Jun 1985-2007/Jul 12
(c) 2007 San Jose Mercury News
File 636:Gale Group Newsletter DB(TM) 1987-2007/Jul 13
(c) 2007 The Gale Group
File 810:Business Wire 1986-1999/Feb 28
(c) 1999 Business Wire
File 813:PR Newswire 1987-1999/Apr 30
(c) 1999 PR Newswire Association Inc
File 2:INSPEC 1898-2007/Jun W4
(c) 2007 Institution of Electrical Engineers
File 35:Dissertation Abs Online 1861-2007/Jun
(c) 2007 ProQuest Info&Learning
File 65:Inside Conferences 1993-2007/Jul 13
(c) 2007 BLDSC all rts. reserv.
File 99:Wilson Appl. Sci & Tech Abs 1983-2007/Jun
(c) 2007 The HW Wilson Co.
File 256:TecInfoSource 82-2007/Jul
(c) 2007 Info.Sources Inc
File 474:New York Times Abs 1969-2007/Jul 14
(c) 2007 The New York Times
File 475:Wall Street Journal Abs 1973-2007/Jul 14
(c) 2007 The New York Times
File 583:Gale Group Globalbase(TM) 1986-2002/Dec 13
(c) 2002 The Gale Group

? ds

Set	Items	Description
S1	2358	DELAY??? (5N) LOCK??? (5N) LOOP???
S2	382	DELAY??? (5N) LOCK??? (5N) CIRCUIT???
S3	2426	S1 OR S2
S4	1762098	MODE OR MODES
S5	13628003	POWER???
S6	66671	S5 (5N) S4
S7	982873	STANDBY OR SLEEP??? OR INACTIVE
S8	25682	S7 (5N) S4
S9	3556802	ACTIVE? ?

14-Jul-07

S10	15304	S9 (5N) S4
S11	94633	S6 OR S8 OR S10
S12	57	S3 AND S11
S13	899210	CLOCK???
S14	2144805	MEMORY? ? OR MEMORIES
S15	25	S12 AND S13 AND S14
S16	6	S15 NOT PY>1997
S17	6	RD (unique items)

read

? t s17/medium,k/all

17/K/1 (Item 1 from file: 16)
DIALOG(R)File 16:Gale Group PROMT(R)
(c) 2007 The Gale Group. All rts. reserv.

07354514 Supplier Number: 58372465 (USE FORMAT 7 FOR FULLTEXT)

Product Times.

Electronics Times, p41
Nov 10, 1997
Language: English Record Type: Fulltext
Document Type: Magazine/Journal; Trade
Word Count: 5793

... has launched two power ranges for its ZX series of single and multiple output switch **mode power** supplies of 300 and 750W.
CE-marked for LVD, the series can be supplied in...

...32bit repeat and instruction word commands which permit data to be input to main system **memory** faster than with a typical direct **memory** access transfer.

Analog Devices
Tel: 01932 266000
Fax: 01932 247401
Enquiry Number 504
Ku-band...C&CD
Tel: 01494 882848
Fax: 01494 882792
Enquiry Number 514

3.3V zero-delay **clock** buffers come in small packages

Cypress Semiconductor has introduced a family of zero- **delay clock** buffers based on phase- **locked loops** (PLLs). The nine devices offer flexible 3.3V **clock** buffering, generating multiple copies of an input frequency with no propagation delay.

They are offered in the industry's smallest packages for zero-delay **clock** buffers -- 8 and 16pin soic and 16pin TSSOP -- and are suitable for **clock** buffering in networking applications such as gigabit and fast Ethernet and ATM, SDRAM buffering in PCs and dimms and **clock** recovery in PC docking stations.

Each device offers a slightly different function. Users can multiply the reference **clock** by two or by four, and can also divide by two or four. The CY2308...migratory 4bit markets.

The TMP87C405M, 408/808M, 409/809M and 408/808L devices all have **clock** gear for low-power operation, and the 28pin package offers considerable space and cost savings...

17/K/2 (Item 1 from file: 148)
DIALOG(R)File 148:Gale Group Trade & Industry DB
(c)2007 The Gale Group. All rts. reserv.

10010106 SUPPLIER NUMBER: 20223991 (USE FORMAT 7 OR 9 FOR FULL TEXT)
Advanced DRAM architectures overcome data bandwidth limits. (dynamic random access memory) (includes related article on the impact of advanced DRAM technologies) (Digital Design)

Bursky, Dave
Electronic Design, v45, n25, p73(9)
Nov 17, 1997

ISSN: 0013-4872 LANGUAGE: English RECORD TYPE: Fulltext; Abstract
WORD COUNT: 6459 LINE COUNT: 00497

Advanced DRAM architectures overcome data bandwidth limits. (dynamic random access memory) (includes related article on the impact of advanced DRAM technologies) (Digital Design)

ABSTRACT: Several semiconductor companies have developed new dynamic random access **memories** (DRAMs) using synchronous data-transfer schemes, byte-serial transfers with dual-edge **clocking** and multibank architectures. Among these new generation of DRAMs are the synchronous DRAM, Rambus DRAM, SyncLink DRAM, Multibank DRAM and cache-enhanced synchronous DRAM. These DRAMs have more **memory** capacity and control options as well as data-transfer rates of up to 1.6...

... architectures to their limits. They also are looking for higher-performance devices that will enable **memory** subsystem performance levels two to four times faster than fast-page mode or extended-data-out DRAMs. **Memory** designers are crafting new generations of dynamic **memories** using synchronous data-transfer schemes, byte-serial transfers with dual-edge **clocking**, and multibank architectures. These new **memory** types will allow designers to build **memory** subsystems that can transfer data at rates up to 1.6 Gbytes/s, allowing the...

...pace with the fastest CPUs and graphics engines.

As data rates go up, so must **memory** capacity - at the high data rates projected for new **memory** chips, current 16- and 64-Mbit devices would be emptied in the blink of an...

...4-Mword by 32-bit chip. Two such chips could provide a 32-Mbyte base **memory** for desktop or portable computers.

There are five architectures competing for the high-speed sockets used in main **memory** subsystems: the synchronous DRAM (SDRAM), which includes its second-generation brother the SDRAM II, and...

...that designers must deal with (see the table).

Graphics subsystems can take advantage of these **memory** types to replace the dual-ported video RAMs that are now used only in legacy graphics subsystems. In addition to the previously mentioned **memory** types, there are some graphics-specific DRAMs that also have been developed - the synchronous graphics RAM (SGRAM), as well as specialty **memories** such as the Window DRAM. However, they will not be covered in this report.

For main **memory** subsystems, there is no clear-cut winner. And in many cases the market will support...

...of megabytes of DRAM, will most likely employ some type of SDRAM for the main **memory**, while the typical future home-office desktop computer will probably initially commit to use some...

...of the RDRAM due to the smaller granularity. In either case, designers are incorporating more **memory** into the systems they create and as the content increases, so will the word width of the **memory** chips to better deal with **memory** system granularity (ILLUSTRATION FOR FIGURE 1 OMITTED).

The issue of granularity is one that continually raises its head, especially as **memory** capacity per chip increases. The first sign of the granularity issue emerged years ago, and...

...designs along with higher densities then made practical 8-, 16- and even 32-bit-wide **memory** chip organizations, and with that the ability to set the desired system granularity. Large **memory** systems would most often employ narrower word widths since such systems often require a lot of depth, while smaller systems often desire wider word chips since the **memory** depth is smaller and wide **memories** could greatly reduce chip

count.

For instance, current 64-Mbit SDRAMs are available with word widths of 4, 8, or 16 bits. Assuming a 64-bit-wide **memory** module, if the unit is assembled with 4-bit-wide SDRAMs, it would have a...

...a depth of 8 Mwords. DIMM versions could double that storage by doubling up the **memory**. And the even larger DIMM modules, such as those used in workstations or servers, could...

...Such storage capacities are so large that small-system users would be hardpressed to afford **memory** upgrades. Graphics subsystems are looking at even wider DRAMs - chips with 32-bit data ports are a better fit since most graphics systems require 8 Mbytes or less of **memory**, and just two to four chips would supply the entire **memory** space.

Wider organizations such as a 4-Mword by 16-bit device would reduce the...

...32 Mbytes using just four SDRAMs. That would permit manufacturers to offer affordable modules. Alternatively, **memories** such as the 16-Mbit RDRAM use a byte-serial approach to transfer data bytes between the **memory** and the system. The host system would then reassemble the wide data words for storage in the cache. Thus, a small **memory** upgrade module that contains as little as 2 Mbytes (one chip) can provide the desired...

...the Nintendo 64 video game, which uses a single 16-Mbit RDRAM for the internal **memory**, a plug-in **memory** expansion module contains just a single 16-Mbit RDRAM, thus doubling the system's **memory** (see "The impact of advanced DRAM technologies," p. 80).

Of all the interfaces discussed and...

...in production. Engineering samples of the cache-enhanced synchronous DRAM are expected shortly from Enhanced **Memory** Systems, and the first prototype of the ...are in design but targeted for sampling in mid-to-late 1998.

Each of these **memory** interfaces has its pros and cons that technically fall into several categories - latency, bandwidth (peak...

...them transfer as little as a page or up to the entire contents of the **memory** chip. The **memory** latency - the time to the first access - is as important as the time for each subsequent data word in the transfer sequence (the burst length).

The sustained bandwidth that the **memory** can achieve is then the burst size (the number of sequential accesses) divided by the...

...terms of scalability, the SDRAM, DDR SDRAM, and SLDRAM can be used in typical 2D **memory** arrays that can expand the word depth or word width and provide simple **memory** subsystem extensions. The RDRAM and its forthcoming Direct RDRAM cousin require a different approach since...

...employ a host-system interface controller that receives the byte/word serial data from the **memory** chips and reassembles the 32- or 64-bit wide-word data or instructions, or breaks wide word information into the byte/word-serial stream for transfer to the **memory** subsystem.

In many systems, a single RDRAM interface controller will typically be used to access...

...sent over the common RDRAM bus. Although the bus is narrower than the wide-word **memory** bus formed by an array of SDRAMs, the higher **clocking** rate (over 600 MHz) gives the bus its high performance. However, one interface controller is...

...needed to increase bandwidth, a second controller would have to be added along with the **memories**, rather than just adding another row of **memory** chips.

In main **memory** systems, latency was very important when cache sizes were small. However with substantial on-chip caches and external level-2 caches, most of the **memory** latency for cache fills is hidden. But on cache misses, the time to reach the...

...critical since without data, the CPU could stall, slowing the entire system. Most of the **clocked memories** do have a two- or three-cycle penalty to reach the first data word, but...

...basically detracted from the device performance when the chips were evaluated for use in main **memory** systems. However, for graphics applications that used long strings of data transfers, the RDRAM delivered ...

...performance and has now found a home on several graphics and multimedia cards.

Souped-Up **Memory**

Many timing and feature improvements have been made to the RDRAM in its second-generation implementation, the Concurrent RDRAM, to optimize it for use in main **memory** subsystems. The Concurrent RDRAM will be available in 8-bit- or 9-bit-wide versions...

...Mbytes/s - a speed faster than any other DRAM approach. Although Rambus has architected the **memories** in conjunction with various partners, it does not manufacture or market the chips. Rather, Rambus licenses the controller interface cell and the **memory** design to its silicon partners. Those licensees include Hitachi, LG Semicon, NEC, Oki, Samsung, and...

...16/18-Mbit chip) row sense amplifier cache, much like the cache DRAM from Enhanced **Memory** Systems. The multiple banks are tied into on-chip control logic that handles the read...

...bank operations yield a high effective bandwidth by using interleaved transactions. For graphics applications the **memories** also include write-per-bit and mask-per-bit capabilities.

Further enhancements to the RDRAM with Intel to better match the **memory** to Pentium CPUs and deliver roughly three times the effective bandwidth of a 100-MHz...

...data path and an 8-bit control bus, with the interface able to operate at **clock** rates of up to 800 MHz (rising and falling edges of a 400-MHz **clock**). DRDRAMs will be available in both 64- and 128-Mbit densities for main **memory** applications, and a 32-Mbit device for graphics applications also is planned for release in late 1998. The chips are essentially the first super-pipelined **memories** that offer a multiple-transaction pipeline and conflict-free transaction interleaving to achieve a sustained...

...electrical interface will be the same as that used in the Concurrent Rambus interface - differential **clocks** and a 1.8-V termination voltage. Internally, the **memory** will use a 128-bit-wide core so that a quad-word (16-byte) transfer...

...8-V signal swing around a 1.4-V reference level, and incorporate some low- **power modes** to better suit them for the mobile systems market (just 0.75 mW on power...

...nap mode, and between 500 and 600 mW when doing reads or writes). Complementing the **memory** chips will be the Rambus access controller (RAC) which can support a single DRDRAM interface or two Concurrent RDRAM interfaces.

Another aspect defined by Rambus for all its **memories** was the packaging, but not just at the chip level. Board layouts and proprietary module layouts also were defined to ensure the **memories** would deliver their best performance. With the DRDRAM, Rambus and Intel have opted to use ...

...support serial presence-detect signals for module identification.

Reducing The Latency

One of the first **memories** to tackle the latency issue - the cache DRAM developed by Enhanced **Memory** Systems - combined a DRAM and a 4-kbit cache on the same chip. The cache is actually formed by the sense-amplifier array and some additional logic and allows the **memory** chip to provide standard access times for a random read (with no matching information held ...

...subsequent read accesses to the same page are read from the cache.)

Although the initial **memory** was only a 4-Mbit device, the company is now sampling a 16-Mbit version...

...FIGURE 3 OMITTED). When fitted with LV TTL I/O (TABULAR DATA OMITTED) buffers, the **memory** will be able to operate at 133 MHz and perform 1-1-1-1 transfers...

...synchronous DRAMs, with their dual-bank architecture were the first mainstream (multiple-sourced) new architecture **memories** to offer performance levels well above that achievable with extended-data-out DRAMs. However, first...

...specified as 100-MHz devices, actually delivered performance levels well below the claimed 100-MHz **clock** speeds.

Although the SDRAMs are designed to a JEDEC standard, slight differences in...

...66 MHz. Subtle spec differences often prevent "blind" interchangeability between different company's ostensibly identical **memories**.

In fact, according to Art Kilmer, manager of **memory** applications at IBM, the latest x86 motherboard chip sets that support SDRAMs have brought out...

...while to figure out which tests are best to assure interchangeability. One issue with the **memories** is the number of banks that can stay open when the host system performs bank...

...the timing and loading issues have yet to be fully analyzed.

To boost SDRAM performance, **memory** designers have tightened some of the ac timing margins, dc parameters, driver characteristics, and layout...

...system operation. (Details are available on the Intel web site: <http://www.intel.com/developer>.) **Memory** designers also are pushing process technology to achieve even higher speeds - **clock** rates of up to 143 MHz will be sampled in early 1998. Specifications for DIMMs...comes to portable systems is that their power consumption is higher than the previous-generation **memories** that used the extended-data-out timing. To overcome that problem, designers at IBM are working on a low-power SDRAM that uses the **clock** enable line to help reduce standby and self-refresh currents by 50 to 75%, while...

...allows the chips to deliver data twice as fast by using both rising and falling **clock** edges, much like the Rambus **memories** use both **clock** edges. The DDR technology is a capability that can be added to any device, thus...

...to meet different market needs. The higher data-transfer rates possible with DDR allow such **memories** to satisfy bandwidth requirements for high-end desktop PCs, high-performance graphics adapters, and workstation servers.

Internally, DDR **memories** are similar to standard SDRAMs, but employ, in the case of the Samsung bidirectional DDR SDRAM, four 512-kword by 32-bit internal **memory** banks, which feed into an output data buffer (ILLUSTRATION FOR FIGURE 4 OMITTED). Additional circuitry to handle strobe generation and timing synchronization (a strobe generator and a **delay - locked loop**) end up increasing the chip's area by about 4% over that of the standard...

...for the DDR chips. The DDR chips also will use differential rather than single-ended **clocking**; however, the DDR parts will give up the **clock** suspend and burst-read/single-bit write capabilities of single-data-rate devices. In terms of ac timing requirements to achieve a 200 Mbit/s data rate from the **memory** pins, that requires a **clock** cycle of 10 ns, an input setup time of 2 ns, an input hold time of 1 ns, a DQS to **clock** time of 1 ns, a data output to **clock** time of 1 ns, and a DQS to data output time of just 0.75 ns.

DDR **memories** actually will come in two variants - one that employs a unidirectional strobe, in which the Read Strobe signal is synchronous with the **clock** signal; and the other version employs a bidirectional strobe that does not have to be synchronous with the **clock**, but rather the data strobe signal is generated by the **memory** controller. The typical read and write timing for a unidirectional DDR **memory** uses the **clock** to latch control and address into the **memory** and a data strobe to latch data into the **memory** controller during the read cycle. During a write cycle, there is no data strobe and all signals and data are referenced to the **clock**.

Unidirectional parts, according to Bob Eminian, manager of DRAM marketing for Samsung, will require a few more pins since the chips will include the **clock** drivers. **Clock** and data lines on all DDR variants, though, will employ series-stub-terminated logic (SSTL...

...have LV TTL interfaces.

In the bidirectional part, the DQS signal is generated by the **memory** controller and is not necessarily synchronous with the **clock**. That provides a larger valid-data window and will permit expandability on unbuffered DIMMs of...

...and are not well suited for implementation in a 4-bit-wide organization for deep **memories**. Use of external **clocks** also limits command-bus bandwidth.

Both implementations will deliver similar system performance, but some systems...

...likely be housed in 100-lead quad-sided flat packages. Such parts would operate at **clock** rates from 125 to 143 MHz. Finally, for the workstation/server type markets expected in...capacities of 256 Mbits with either a 4- or 8-bit-wide data word and **clock** at 125 to 143 MHz. Such devices would carry no legacy-compatibility requirements and come...

...a new package to deal with the more stringent timing requirements.

One key issue as **memory** chips hit capacities of 64 Mbits and higher, is packaging. Designers at Rambus were the first to acknowledge the need

for special packages to deal with high-speed **memories** and crafted unique vertical incline packages that reduced lead lengths to a bare minimum. Today, chip-scale packaging and ball-grid array technology are being deployed, but the bulk of **memory** chips are still targeted to go into SO or QFP style packages.

In the SO...

...forthcoming 256-Mbit devices in a 400-lead-wide package, saving about 47% of the **memory** board space as would have been required with traditional first-generation 500-lead packaged **memories**. That also should eliminate the need for a major module redesign when moving from the 500-lead 64-Mbit devices.

An alternative to that is a stacked **memory** approach that IBM and other companies are pursuing. Reminiscent of the early attempts in the...

...and stacking two or four of them on top of each other to form tiny **memory** modules that can be mounted on DIMM PC boards. This approach can produce modules four...

...be offered in both a 64-lead vertical surface-mount package that allows high-density **memory** arrays to be constructed, and in an 80-lead TSOP. Initially expected to be implemented...

...SLDRAM? Well, for starters, it looks very similar to an SDRAM, but packs eight internal **memory** banks, a **clocked** synchronous interface that is terminated with small-swing signaling, and employs a programmable data burst...

...approach decouple the internal DRAM address and control paths from the data interface, allowing the **memory** to achieve higher bandwidth data transfers. Like DDR SDRAMs, SLDRAMs will use both rising and falling **clock** edges to transfer data, and a return **clock** signal to help improve timing margins.

Internally, an SLDRAM consists of a **memory** array organized as eight banks, each structured as 128 kwords by 72 bits, and each...

...and ID assignment are just some of the setups that must be done. A SyncLink **memory** chip may have multiple subRAMs or blocks, and except for initialization, which is at least...

...blocks act essentially like independent RAMs and each handles one request at a time.

The **memory** chips are connected by commands and data links - the controller drives the command link to send Read, Write, Load, Store and Event commands to the **memories**. The datalink is driven by an SLDRAM and received by the controller in the case...

...bit command link bus, a two-byte datalink bus, and Select, LinkOn, and other control/ **clock** signals, which connect to an array of SLDRAM that shares common command and data buses...

...power-up to synchronize the SLDRAMs and assign unique IDs to each.

One other novel **memory** type is available to provide higher **memory** system bandwidth - the multibank MoSys DRAM, which provides a high-bandwidth 16-bit interface that can transfer data at up to 666 Mbytes/s when **clocked** at 166 MHz. Internally, MDRAMs use 32 banks per megabyte. That will permit the chips...

...MDRAMs is that they require short buses and that, in turn, limits the number of **memory** chips to just four. Current density levels of the parts range from 0.5 to...

...per chip, and even higher densities are on the drawing boards. Thus, although for limited **memory** applications the MDRAMs can serve as main **memory**, they would appear to be better suited for high-speed buffers and graphics applications.

References:

B. Prince, High-Performance **Memories**, Wiley, 1996

S. Przybylski, New DRAM Technologies: A Comprehensive Analysis Of The New Architectures, MicroDesign Resources, 1996

Companies Mentioned In This Report

Enhanced **Memory** Systems Inc. 1850 Ramtron Dr. Colorado Springs, CO 80921 (719) 481-7003 <http://www.edram...>

...terms of new chip sets, tighter timing specs for both SDRAM chips and SDRAM-based **memory** modules, plus a higher awareness on the part of both PC OEMs and the supplier...

...architecture seems to have been problematic, the technical difficulties of transitioning to even more advanced **memory** architectures are incrementally more severe.

Workstations and PCs are pointed toward an abruptly upscaled performance...

...new systems will require synchronization and tight coupling between the CPU, the chip set, the **memory** modules, I/O ports, motherboard, and peripherals far beyond anything that's been achieved so far.

Even in the most advanced systems available, with the SDRAM **memory** bus operating at 100 MHz, there is only 10 ns available for a **clock** cycle. The SDRAM itself requires up 6 ns, leaving about 4 ns for all other ...

...the front runner, are promising operating frequencies of 400 to 600 MHz, which will drive **clock** cycles down to 2.5 and 1.6 ns, respectively.

For independent **memory** -module manufacturers, continues Johnston, the design constraints for these advanced modules will contribute to a require 660-MHz testing to adequately test the parts. Few independent **memory** -module suppliers are positioned today to provide even the 330-MHz capability and have the...

...equivalent with competing DRAM architectures such as SDRAM or DDR DRAM, raises the ante for **memory** module design competence. There are new design challenges as systems transition from simple connection of...

...microstrip-line and transmission-line design.

To simulate signal-integrity of the module design, the **memory** controller model, the connector model, and of course the detailed modeling of the traces and...

...can be included in their system simulations. In addition to these design constraints, if any **memory** -module maker is not, for example, already implementing microball-grid-array (MBGA) packaging competence, it...

...chance of surviving this next transition phase.

The implications also are severe for buyers of **memory** modules, whether they are PC OEMs or individuals in the aftermarket. It is likely that...

...system will meet performance requirements. PC OEMs will partner much more tightly with carefully selected **memory** module suppliers, and the **memory** aftermarket may be controlled exclusively by the OEMs who

manufacture the systems, with consumers buying...

DESCRIPTORS: Dynamic random access **memory** --...

... **Memory** (Computers

17/K/3 (Item 2 from file: 148)

DIALOG(R)File 148:Gale Group Trade & Industry DB
(c)2007 The Gale Group. All rts. reserv.

07714326 SUPPLIER NUMBER: 16682578 (USE FORMAT 7 OR 9 FOR FULL TEXT)
**Gigabit DRAMs, 64-bit CPUs and more at ISSCC. (dynamic random access
memory ; central processing unit; International Solid State Circuits
Conference) (includes related article)**

Burksy, Dave

Electronic Design, v43, n4, p61(13)

Feb 20, 1995

ISSN: 0013-4872

LANGUAGE: ENGLISH

RECORD TYPE: FULLTEXT; ABSTRACT

WORD COUNT: 9237 LINE COUNT: 00710

**Gigabit DRAMs, 64-bit CPUs and more at ISSCC. (dynamic random access
memory ; central processing unit; International Solid State Circuits
Conference) (includes related article)**

...ABSTRACT: at the International Solid State Circuits Conference held in
Feb. 1995. Two dynamic random access **memories** (DRAM) with gigabit
complexities, advanced 64-bit microprocessors and several digital
technologies were showcased during...

... with gigabit complexities, and also highlighted were developments
in other high-performance DRAM and flash- **memory** areas, as well as
advanced 64-bit microprocessors and many other digital technologies.
Although years...

...incorporated to allow the systems to operate faster. This applies to
other digital ICs besides **memories** and microprocessors (see "There's
digital life beyond **memories** and microprocessors," p. 66). These
considerations can be seen in experimental 1-Gbit and 256...

...a random-access approach, and uses synchronous timing for high
data-transfer rates. Additionally, the **memory** incorporates an output
buffer that cancels ringing, ensuring reliable high-speed data transfers.
Large chip...

...of metal are employed. Along with the large chip area, however, comes
the problem of **clock** distribution. To overcome this obstacle, Hitachi
developed a distributed-column-control architecture that permits
independent operation of the I/O circuitry and the **memory** subarrays
[ILLUSTRATION FOR FIGURE 1 OMITTED].

A single external **clock** signal is supplied only to the I/O buffers.
Each of the 32-Mbit subarrays is controlled by a local timing generator,
driven by the externally supplied **clock** signal and transitions on the
address bus. A FIFO buffer was added to the I...

...buffer circuits to help compensate for timing differences between the
I/O buffers and the **memory** array.

Organized as 64 Mwords by 16 bits, the DRAM operates from 1.5 V with a
row-address-strobe (RAS) access time of 33 ns (six 4.5-ns **clock** cycles
plus a 6-ns setup time). A burst transfer mode incorporated into the
synchronous...

...approach in regards to their 1-Gbit DRAM for large-file-storage applications. This file **memory** performs word serial data transfers at a high rate - 400 Mbytes/s over a 32...

...I/O path - thanks to a pipelining scheme and the use of a 100-MHz **clock**. The first access requires 15 cycles of the 100-MHz **clock** (about 150 ns) to position the **memory** pointers. The chip consumes relatively little power - just 68 mA when operated from 2V - because...

...issue, NEC implemented a time-shared offset-cancelling sensing scheme and a diagonal-bit-line **memory** cell. This combined approach cuts the chip's area by 30% when compared to a...the operating voltage to 1.2 V while accessing data in 49 ns. In the **standby mode**, current drops to a mere 5 [[micro]ampere].

A charge-transferred, well-sensing scheme is...

...V_{wp}) are set to V_{cc}/2 and ground. When data is first read from the **memory** cell, the voltages on SN and V[W.sub.p] are shunted in response to ...

...can control the local power-line levels in response to moving between the normal or **sleep modes**.

Designers at the Semiconductor Research Center at Matsushita Electric Industrial Co. Ltd., Osaka, Japan, have...

...two rows plus 16 columns of redundancy per 512-kbit block.

Wave pipelining allows the **memory** to operate at 150 MHz, and offers a programmable column-address-strobe (CAS) latency. Divided into eight 32-Mbit **memory** banks and organized as separate 2-Mword-by-16-bit subarrays, the **memory** chip also has a hierarchical I/O architecture with evenly distributed line capacitances, thereby producing well-defined delays.

The internal **clock** buffer uses a 150-MHz input **clock** to generate short synchronization pulses. The pulses are created by post-charge logic, a novel...

...The feedback signal has a finite delay, which determines the pulse width of the internal **clock** signal. The post-charge logic was added throughout the **memory** chip wherever a sharp pulse edge was needed. Hyundai estimates that the use of post-charge logic reduces the access time in the typical **memory** data path by about 10%.

Wave pipelining was applied to the column-access path. However...

...data pulse into the pipelined data path, without waiting until the previous pulse has been **clocked** into the receiving data-out-put-buffer register.

Clocked storage elements are not included the column-access path, thus both **clocking** and partitioning overheads can be eliminated. This allows wave pipelining to increase the **clock** rate while at the same time the total latency time remains practically unchanged.

Total delay...

...is just 18 ns, and the data pulse width is less than 6 ns. The **clock** frequency then can be raised up to 150 MHz without the insertion of any additional **clocked** latches.

3D GRAPHICS **MEMORY**

Graphics subsystems can readily take advantage of the high speeds possible with synchronous, wave-pipelined **memories**. General-purpose DRAMs, however, aren't usually optimized for 3D-graphics frame buffers,

since they...

...Z-compare and Alpha-blending. As a result, Mitsubishi developed a 3D graphics frame-buffer **memory** in conjunction with Sun Microsystems Computer Corp., Mountain View, Calif. In the meantime, Toshiba Corp...

...graphics and includes both a Z-compare and an Alpha-blend unit along with the **memory**. It can greatly accelerate many graphics operations, delivering rendered image data at up to 400...

...The latch de-couples row-access and column-access cycles from one another, permitting the **memory** cell access to be performed behind the data-latch (column) access. Random accesses take place in just 10 ns when the array is **clocked** at 100 MHz.

FASTER SRAMs

Although the high speeds of special-architecture DRAMs have caught...

...wave pipelining for 300-MHz operation, Hitachi's SRAM starts with a 150-MHz external **clock**. It generates 300-MHz multiphase pulses by using a multiphase phase-locked loop that controls...design rules and four levels of metal, the macro consists of 128 identical 8-kbit **memory** blocks, each of which is complete, aside from address decoders. A simple change in the ...

...Ltd., Kawasaki, Japan, and HaL Computer Systems Inc., Campbell, Calif., also yielded a high-performance **memory** macro that HaL incorporated into a 64-bit processor. The macro is a 14-port **memory** used in the read-renaming register file. **Memory** consists of 116 words, is 64-bits wide, and has ten read and four write...

...fast context switching. Implemented as a multichip module (MCM), it consists of the CPU, a **memory**-management unit (MMU), four cache-**memory** chips, and one **clock** chip, for a total of almost 22 million transistors. The processor subsystem can operate at 154 MHz [ILLUSTRATION FOR FIGURE 7 OMITTED].

The four 64-kbyte cache **memories** are evenly divided to form 128-kbyte data and instruction caches, while the MMU provides a 128-bit-wide interface to main **memory**. The MCM approach, however, allows even wider buses to be used if higher operating bandwidth...

...The second-level caches are virtually indexed and tagged. Each of the specially designed cache **memory** chips is four-way set associative and can service two independent requests from the CPU. Latency from address generation to data use is just three **clock** cycles. The MMU, responsible for managing the **memory**, maintaining **memory** coherence, and error handling, has three levels of address spaces: (1) a virtual space for...

...space for I/O devices and the diagnostic processor, and (3) a physical space for **memory**. These hierarchical spaces provide an efficient mechanism for managing the large, 64-bit address space...integer, floating-point, and flag variables. Also incorporated are multiprocessing bus support and carefully controlled **memory**-access reordering.

The engine predicts the instruction flow, and the instructions are decoded into micro...

...load and one store per cycle. The L2 cache interface runs at the full CPU **clock** speed and can transfer 64 bits every bus cycle. The bus can operate at 1/2, 1/3, or 1/4 the CPU **clock** speed, which Intel designers expect will yield processors with throughputs of 200 SPECint92.

Nexgen Inc...

...for cost-sensitive consumer applications. The PowerPC 620, which can operate at a 133-MHz **clock**, issues four instructions every **clock** cycle for a throughput of 225 SPECint92 and 300 SPECfp92, when used with a 4...

...caches feed the processor's pipelines. Associativity is achieved through the use of content-addressable **memories** (CAMs) within binarily decoded sections of the cache. This is in contrast to the usual...

...the use of more-complex CAM cells.

To speed along instruction execution, the 620's **memory**-management subsystem employs a unique two-level translation scheme that allows small cycle times as consuming just 300 mW when **clocked** at 50 MHz, the processor achieves a 150-MIPS/W rating by delivering a throughput...

...software breakpoint interrupts, single-step execution, real-time program-counter tracking, data-cache locking, and **power**-down **modes**. To minimize **active** on-chip **power**, a 128-bit instruction queue was added to reduce the amount of time the instruction...

...bit adiabatic pulsed-power-supply (APPS) multiplier that trims power to a bare minimum. When **clocked** at 13.9 MHz, the multiplier consumed only about 20 mW. At the same speed...

...the ramp, the chip is powered down to [V.sub.SS]. The states of all **memory** elements are stored on the parasitic capacitances and can be held for 1 ms without...

...The chip operates from 1.5 V and consumes 1 W. For high area and **power** efficiency, multiple-valued current-**mode** MOS differential logic circuits are used.

A special threshold detector was developed based on differential...

...portable, low-power systems, is now leveraging flash-storage schemes with NOR, NAND, and AND **memory** architectures to achieve densities of 32 Mbits and beyond. Four ISSCC papers discussed 32-Mbit electrically reprogrammable **memories**, while two additional presentations focused on 16-Mbit devices.

Leveraging its 32-Mbit flash technology...

...described a scheme for doubling the capacity of the chip without doubling the number of **memory** cells. A four-level threshold detector is incorporated into the sense amplifiers. Instead of a...

...10, and 11-bit patterns, providing enough states to define two independent bits in each **memory** cell (ELECTRONIC DESIGN, Sept. 19, 1994, p. 46). Special read-reference cells (R1, R2 and...

...binary search sensing scheme (BSSS) and are compared to the measured threshold voltage in the **memory** cells to determine the cell value.

Cell programming is done through the controlled application of...time of 120 ns.

Fast programming time is a long sought-after goal with flash **memories**. The faster the programming, the more the chips can be used in main-storage applications...

...current 10-to-100-fold speed difference between the read and store functions on flash **memories** is slowly being whittled down, as several companies demonstrated. For instance, a 512-byte on...

...other companies, is a good match for file storage and archiving. The

AND-based flash **memories** inject electrons onto the floating gate through Fowler-Nordheim tunneling to raise the threshold voltage...

...programming signal in accordance with the repetition of verification, since the threshold voltage of the **memory** cells depends logarithmically on the program time. This approach reduces the time overhead for verification...

...technique is a key aspect of a 3.3-V, 32-Mbit, NAND-based flash **memory** described by Samsung Electronics Co. Ltd., Kiheung, Korea. The ISPP scheme allows fast programming of the **memory** (2.3 Mbytes/s) by using an erase-block size of 8 kbytes (16 pages...

...blanked in a short 7 ms. Interleaved data paths on the chip also allow the **memory** to achieve a read throughput of 24 Mbytes/s on the 4-Mword-by-8-bit circuit.

In the NAND **memory** structure, a unit NAND string consists of 16 **memory** cells and two select transistors. To minimize the die size, blocks of cells are organized...

...page-size units of 512 bytes.

Another 3.3-V, 32-Mbit, NAND-based flash **memory**, described by Toshiba, achieves a 35-ns cycle time for data reads and programming data...

...strapped select gates and boosted word lines to reduce the read-out access time. The **memory** also incorporates special erase-block registers that allow users to specify more than one block for simultaneous erasure.

The **memory** chip divides the program and read operations into two parts. The first part is data transfer between the **memory** cell and data register connected to each bit line. This transfer is performed by simultaneous...

...Mbyte/s. The CMOS chip employs Mitsubishi's proprietary divided bit-line NOR (DINOR) flash-**memory** architecture, and uses a 256-byte page buffer. Made with 0.5-[[micro]meter] design...

...programming of the 64-kbyte subblocks that make up the 2-Mword-by-8-bit **memory** chip. Selective blocks can be locked to prevent erasure. The controller also has an erase...

...was achieved by adding a new interface that provides synchronous signal timing. The synchronous flash **memory** is the first to employ a **clocked**, synchronous interface that allows full random access to the contents, and continuous burst transfers at...timing-relationship information. In addition, an on-chip programmable-internal-latency (PIL) register tells the **memory** the number of cycles it will need to complete an access in each bank. For...

...indicates that a two-cycle delay should be added to the words delivered to the **memory** interface. The interface then matches its data delivery timing to any subsystem.

The access is...

...valid line and supplies valid addresses on the next four rising edges of the external **clock**, T0 to T3. The addresses are sequential, and are thus loaded into alternate banks, getting...

...data from the request at T0 is latched and driven out of the chip two **clocks** later at T2, the data from the request at T1 is latched and driven out at T3, and so on.

FERROELECTRIC **MEMORIES**

Although oxide has been the most common storage dielectric material used for nonvolatile electrically programmable **memories**, research over the last few years has focused on ferroelectric **memory** cells. These cells are attractive because they offer much higher endurance levels and are simple to fabricate.

The latest effort in ferroelectrics, a single-transistor ferroelectric **memory** cell, comes from Rohm Co. Ltd., Kyoto, Japan. The novel ferroelectric floating-gate RAM consists...

...individual cell with the FETs lined up in a simple matrix. The use of a **memory** cell consisting of a single transistor - an enhancement-type p-channel FET - solves some of...

...the addition of another element, such as a variable resistor or a diode in each **memory** cell.

RELATED ARTICLE: THERE'S DIGITAL LIFE BEYOND **MEMORIES** AND MICROPROCESSORS

Significant advances in digital areas such as high-speed I/O and **clock** buffers, digital-signal processors and quantum electronics are often overshadowed in the face of highly-profiled developments in **memories** and microprocessors. At ISSCC, for example, Intel unveiled details on a 900-Mbit/s bidirectional...

...lock buffer that automatically compensates for skew across multiple outputs. The buffer has six digital **delay - locked - loop** (DLL) remote **delay** regulators that all share a common **clock** -reference input. Each regulator is dedicated to driving a single external load. Load characteristic updates...

...2.4 to 5 [[micro]seconds] to keep skews under 1 ns maximum from the **clock** chip to the targeted point of use.

Regulator skew compensation starts with a replica loop that models board-etch and **clock** -chip propagation delays. The loop reproduces delays in the **clock** chip I/O lines and combinatorial logic delay overhead, and the loop's output drives...

...load. This concept does, however, require a joint effort. Simply put, each load must have " **Clock** Before" and " **Clock** After" taps that return to the associated regulator in the **clock** buffer chip through equal-length ...circuits between the audio and video functions. Thus, the chip spends about 15% of its **clock** cycles on audio decoding, about 80% on video decoding, and about 5% on system stream...

...s variable-pipeline architecture permits the horizontal search range to be expanded without external logic. **Clocked** at 40 MHz, the chip has a peak computational throughput of about 20 gigaoperations/s...

...implemented with just two vMOS devices and a few switching transistors (standard MOSFETs).

A new **clocking** scheme from Tohoku cancels fluctuations in threshold-voltage shifts of the vMOS transistor that arise from fabrication processes, improving noise margins. The **clocking** approach starts with a **clock** -driven switching transistor that's attached to the floating gate of the vMOS device. That...

17/K/4 (Item 1 from file: 275)
DIALOG(R) File 275:Gale Group Computer DB(TM)
(c) 2007 The Gale Group. All rts. reserv.

01902075 SUPPLIER NUMBER: 17946197 (USE FORMAT 7 OR 9 FOR FULL TEXT)
Overview of code-domain power, timing, and phase measurements. (Technology Tutorial) (Tutorial)
Birgenheier, Raymond A.
Hewlett-Packard Journal, v47, n1, p73(21)
Feb, 1996
DOCUMENT TYPE: Tutorial ISSN: 0018-1153 LANGUAGE: English
RECORD TYPE: Fulltext; Abstract
WORD COUNT: 15152 LINE COUNT: 01281

... tau).sub.0) and (delta)(omega)) are known a priori, then the "Fast Code-Domain **Power** " **mode** should be used.

Presented in Fig. 18 are curves obtained from simulations showing the rms...the HP-UX port to PA-RISC and on the HP 9000 Model T500 processor-**memory** bus definition. He led the protocol definition team and the verification effort for the HP...

...of future systems. His work has resulted in nine patents in architecture, cache design, and **memory** design. He has coauthored three articles on PA-RISC, the HP 9000 Model 840, and...
...bus interface and participated in the Runway bus definition. He is currently working on the **memory** controller design for the next generation of servers. He earned a BSEE degree from the...

...He is married and his hobbies include bicycling and woodworking.
(*) The chip interval is the **clock** period of the spreading code used in a spread-spectrum system. In this paper, a...

...coming to HP, he has been responsible for on-chip circuit design of several processors, **memory** controllers, and bus converters on several HP 3000 and 9000 systems. Most recently, he designed...electronics complex of next-generation systems. His work has resulted in a patent on a **delay - locked loop circuit** and he has coauthored two papers on VLSI processors. Francis is married, has a daughter...

...he worked on the design of the processor interface chip, bus converter chip, and processor-**memory** bus definition for the HP 9000 Model T500 computer system. Before that, he worked on...

...CPU. Before that, he worked on electrical verification for HP 3000 Series 990 and cache **memory** design and on electrical verification of the HP 3000 Series 990. His work has resulted...the HP 9000 J-class workstations and K-class servers. He is currently responsible for **memory** controller control logic design. Jim was born in Ames, Iowa. He is married and is...

...Muzaffarnagar, India, Akshya is married, has two children, and enjoys cricket, tennis, and camping.

44 **Memory** System Design

Thomas R. Hotchkiss

A hardware product planning engineer, Tom Hotchkiss is currently responsible...

...lead engineer and architect for the highperformance VLSI chipset for the HP 9000 K-class **memory** subsystem. Before that, he was a VLSI designer for a pair of HP PA-RISC...

...he worked on the design of the SIMM boards for the HP 9000 K-class **memory** system. Before that, he served as a design engineer or project manager for HP 3000...

...33 computers and did VLSI chip design for the HP 3000 Series 50 and 60 **memory** subsystem controller and for PA-RISC I/O and CPU chips. He contributed to the **memory** subsystem design as well as the board design and layout for the **memory** carrier board used in the HP 9000 J/K-class systems and HP 3000 Series...he is a principal member of the technical staff, currently responsible for 622-Mbit/s **clock** data recovery (CDR) postamplifier design and fibre channel arbitrated loop IC design. For the fiber...

17/K/5 (Item 1 from file: 621)

DIALOG(R)File 621:Gale Group New Prod.Annou.(R)

(c) 2007 The Gale Group. All rts. reserv.

01047179 Supplier Number: 40113035 (USE FORMAT 007 FOR FULLTEXT)

CHIPS INTRODUCES FIRST IBM MODEL 30-COMPATIBLE VLSI CHIPSet

News Release, p1

July 20, 1987

Language: English Record Type: Fulltext

Document Type: Magazine/Journal; Trade

Word Count: 866

... for 8088, 8086, V20 and V30 microprocessors, the 82C100 Model 30-compatible chip includes a **memory** controller that supports the Lotus-Intel-Microsoft expanded **memory** specification (EMS), system configuration registers, and dual **clocks** . It also offers power management features for laptop systems to help reduce battery consumption when...

...Convertible.

The chip includes an 8237-compatible DMA controller, 8259-compatible interrupt controller, 8284-compatible **clock** generator, 8288-compatible bus controller, 8254-compatible timer/counter, 8255-compatible peripheral port, and parity...

...printer and scanner

functions, and a game port. The CHIPSpak also includes a real-time **clock** and a low- **power** serial port. **standby mode** that is particularly useful in laptop computer applications. There are 114 bytes of CMOS SRAM...

...82C605A CHIPSport has the same features as CHIPSpak, but does not include the real time **clock** . It is ideal for use in PC AT systems, add-on boards, and standard system board configurations that already contain a real time **clock** -- either as a discrete component or integrated into the system via the company's Integrated...

...high

performance floppy disk controller subsystem.

The data separator features a self-calibrating analog phase- **locked loop** , write precompensation, and DRQ **delay circuitry** . It also supports multiple data rates of 250K, 300K and 500K bits per second, all...

17/K/6 (Item 1 from file: 810)

DIALOG(R)File 810:Business Wire

(c) 1999 Business Wire . All rts. reserv.

0058141 BW626

CHIPS & TECHNOLOGIES: Chips introduces first IBM

July 20, 1987

Byline: del 30-compatible Vlsi Chipset

...for 8088, 8086, V20 and V30 microprocessors, the 82C100 Model 30-compatible chip includes a **memory** controller that supports the Lotus-Intel-Microsoft expanded **memory** specification (EMS), system configuration registers, and dual **clocks** . It also offers power management features for laptop systems to help reduce battery consumption when...

...convertible.

The chip includes an 8237-compatible DMA controller, 8259-compatible interrupt controller, 8284-compatible **clock** generator, 8288-compatible bus controller, 8254-compatible timer/counter, 8255-compatible peripheral port, and parity...

...printer and scanner functions, and a game port.

The CHIPSpak also includes a real-time **clock** and a low- **power** serial port **standby mode** that is particularly useful in laptop computer applications. There are 114 bytes of CMOS SRAM...

...82C605A CHIPSport has the same features as CHIPSpak, but does not include the real time **clock** . It is ideal for use in PC AT systems, add-on boards, and standard system board configurations that already contain a real time **clock** -- either as a discrete component or integrated into the system via the company's Integrated...

...high performance floppy disk controller subsystem.

The data separator features a self-calibrating analog phase- **locked loop** , write precompensation, and DRQ **delay** **circuitry**

. It also supports multiple data rates of 250K, 300K and 500K bits per second, all...
?